Inertial Regularization and Selection (IRS)—Sequential Regression in High-Dimension and Sparsity

Chitta Ranjan
Georgia Institute of Technology
Atlanta, U.S.A.
nk.chitta.ranjan@gatech.edu

Samaneh Ebrahimi
Georgia Institute of Technology
Atlanta, U.S.A.
samaneh.ebrahimi@gatech.edu

Kamran Paynabar
Georgia Institute of Technology
Atlanta, U.S.A.
kamran.paynabar@gatech.edu

arXiv:1610.07216v1 [stat.ML] 23 Oct 2016

ABSTRACT

In this paper, we develop a new sequential regression modeling approach for data streams. Data streams are commonly found around us, e.g in a retail enterprise sales data is continuously collected every day. A demand forecasting model is an important outcome from the data that needs to be continuously updated with the new incoming data. The main challenge in such modeling arises when there is a) high dimensional and sparsity, b) need for an adaptive use of prior knowledge, and/or c) structural changes in the system. The proposed approach addresses these challenges by incorporating an adaptive L1-penalty and *inertia* terms in the loss function, and thus called Inertial Regularization and Selection (IRS). The former term performs model selection to handle the first challenge while the latter is shown to address the last two challenges. A recursive estimation algorithm is developed, and shown to outperform the commonly used state-space models, such as Kalman Filters, in experimental studies and real data.

CCS Concepts

•**Computing methodologies** → **Machine learning algorithms;**

Keywords

regression; data stream; high-dimensional

1. INTRODUCTION

In recent times, huge amount of data is collected every day for most of statistical learning processes. For eg., Twitter, Facebook, Walmart, Weather, generates tremendous amount of data everyday. This leads to accumulation of large amounts of data over time. Such processes are also sometimes called as a data-streaming process. Modeling an entire datastream together is computationally challenging. Although, methods like Stochastic Gradient Method has been proven to work on such large data, it will fail when the process is also changing over time. We call such changing systems as evolving.

Take for example, the behavior of users on YouTube evolves over time. The user usage data is collected by YouTube, which should be used for continuously modeling and updating their recommender system. Another example can be found at retailers demand forecasting models. For eg., at Walmart several thousands of products are sold. For each type, several new brands are introduced or replaced every day. The demands of all these products interact with each other. The model should, thus, be continuously updated to accomodate for the dynamic changes made in the demand system.

Sequential modeling of such processes can alleviate the challenges posed by the dynamic changes in evolution in the system. Moreover, it significantly reduces the computation complexities. Several researchers have worked in this direction to develop methods specific to certain problems, for eg. evolutionary recommendation systems [6].

Broadly, the existing methods use a data-pruning or sequential models. A data-pruning approach takes the most recent data, either using a moving window or giving decreasing weight based on the oldness of the data. These methods are, however, ad-hoc with the window size or the weight functions difficult to optimize. Moreover, they do not appropriately use the model information learnt so far.

Sequential models, on the other hand, are better at using the past information on the model and updating it from the newly observed data. However, the commonly used sequential methods, like Kalman Filters, Particle Filter, etc. are not suitable for high-dimensional data. Moreover, they are not capable of optimally selecting a sub-model when there is sparsity.

On the other hand, the abovementioned problems are high-dimensional and sparse. Besides, following are the inherent challenges in the problem at hand,

1. *Resemblance*: On updating the model from new data, any abrupt model change can be difficult to interpret. For eg., suppose in the previous model a predictor variable, *age* and *sex* are significant and insignificant, respectively, but after an update in current epoch, the significance of *age* and *sex* reversed. Such abrupt changes can easily occur, due to multicollinearities and often performed suboptimal search of objective space and, thus, many possible solutions. However, interpreting adruptly changing solutions is difficult and undesirable.
2. *Computation*: Since the data size will increase expo-

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

nentially, modeling all the data is computationally impractical. An effective sequential modeling is required to break the computation into smaller parts.

3. *Robust* (to noise and missing values): Some variables can be unobserved (or scantily observed) in certain epochs. For example, at Walmart several products get added/removed everyday or suspended for some time. It is still essential to keep track of these dynamic changes and appropriately handle the model. Also, there can be a large amount of noise occuring in an epoch (due to any external event). For example, someone else using your Netflix for a short while, often leading to your recommendations going of the chart.

4. *Accuracy*: Obtaining high prediction accuracy in a changing and high-dimensional process is difficult. Specifically, overfitting is a common issue in modeling a high-dimensional sparse model with poor model selection, leading to a high test error.

It is, therefore, important to develop a new method that can sequentially model a high-dimensional data streaming process, with inherent sparsity and evolution. To that end, we propose a method called Inertial Regularization and Selection (IRS). Similar to state-space models, IRS sequentially models them over subsequent time epochs. A time epoch can be a day, a week, or any other time interval at which the model is required to be updated.

IRS has an *inertia* component in the objective loss function. This inertia component resists abrupt model change, bringing *resemblance* with model learnt in earlier epochs. This also enables IRS to use the prior knowledge learnt on the model so far, together with the current epoch's data, for a yield an accurate model. It also keeps the model smooth and robust to intermittent noisy observations. Besides, unlike other state-space models, it works better in high-dimensional and sparse problems due to its automatic optimal model selection property. An optimal sparse model estimate prevents overfitting, and thus, lowers test prediction errors.

IRS is a domain agnostic (generic) methodology that can handle any large or small statistical learning problems with different specifications. In this paper, it is developed for linear relationships and gaussian errors, however, it can be easily extended to non-linear scenarios using a link function. We show existing related methods as a special case of IRS, viz. Kalman filters, (adaptive) Lasso and Bayesian models. However, these models have certain critical limitations, especially in terms of ability to handle high-dimensions, use of past information (sequential modeling) and computation, discussed more in §2 below, which is overcome by IRS. Through development of IRS, this paper also shows the connection between Kalman Filters, regularization methods, and Bayesian frameworks. As a by-product, we also show an estimation procedure for Kalman filter which has significantly faster computation in the usual case of data size being larger than the number of predictors.

In the remaining of this paper, we first review the related work in this area in §2. Thereafter, we show the development of IRS in §3. This section provides the definition and elaboration of IRS, followed by development of IRS regularization parameters (§3.3). Besides, within this section, equivalence of IRS with other methods as special cases are shown in §3.4-3.6. Further, we present a closed-form solution under orthogonality conditions (in §3.8.1) and an iterative solution, for any general case, in Sec. 3.8.2. An algorithm based on the iterative solution is postulated in §3.8.3, which also discusses various computational and implementation aspects of IRS. Furthermore, we experimentally validate IRS's efficacy in §4, and show its application on a real online retail data in §5. Finally, we discuss the paper and conclude in §6.

2. RELATED WORK

Sequential modeling of statistical learning, specifically in context of regression, was attempted by Gauss (1777–1885) more than a century ago. However, his method, called as recursive least-squares (RLS) estimation came into light several years later after a rediscovery by [14]. RLS works by estimating regression parameters in a model with the available data and updates it when new data is observed. However, it works if the model parameters are expected to be stable, i.e. the model is not evolving. Moreover, since the entire data until the present is used for the model, the effect of successive new data on model starts to diminish. To account for this, a rolling regression, which a data-pruning approach that keeps only a fixed window of data for modeling. A gradual discounting of old data using some decaying function is also used. Nevertheless, RLS still suffers from its inability to effectively model an evolving process, and also in handle high-dimensional data.

To better account for the model evolution, Kalman filters (KF) were developed[10]. KF is a sequential learning procedure, where the model at any any epoch is estimated as a weighted average of its predicted state from the previous state and the estimate from the new data. For a consistent estimation, KF incorporates the model's covariance for the model reestimations — the weights in the weighted average proportional to this covariance —, but it is not designed to handle high-dimensional data. Other generalizations of KF, viz. extended Kalman filter and unscented Kalman filter [9] were developed for nonlinear systems, but also lacked in their ability to model high-dimensional data.

Ensemble Kalman filter (EnKF)[7], was developed for high-dimensional situations. It is a combination of the sequential Monte Carlo and Kalman filter. However, EnKFs are majorly focused on a proper approximation of the model's covariance matrix in situations of high-dimension but less data (via. Monte Carlo approximations). Here an ensemble a collection of model realizations from different Monte Carlo generated data replications, from which the model covariance is approximated as the sample covariance. As can be noted, it does not address the issue of sparse model selection under high-dimensionality. While some "localization" techniques are some times used for dimension reduction[13], for eg. in geophysics the model is decomposed into many overlapping local patches based on their physical location, filtering is performed on each local patch, and the local updates are combined to get the complete model. However, this technique is not easy to generalize to any model.

Particle filters (PF)[8] are another class of method for state-space modeling under nonlinearity. It is usually seen as an alternative to KFs in nonlinear scenarios. But they severely suffer from the "curse of dimensionality" due to its nonparametric nature even for moderately large dimension[12]. [11] developed an ensemble filtering method by combining the EnKF and PF for high-dimensional settings.

They used a similar "localization" technique described above, which limits its application.

State-space models, specifically Kalman filter, have been extensively developed in the field of signal processing. To handle high-dimension and sparsities in signals the methods extended the Kalman filters to address them. [17] proposed a Kalman filtered compressed sensing, which was further extended in [3], who directly modified the Kalman filter equations using compressed sensing and quasi-norm constraints to account for sparsity. However, these methods lack theoretical support and has limited application. [4, 5] showed different signal regularization possibilities for different signal dynamics scenarios. They also showed an extension of [1] and [2] by adding an L_2 penalty in the optimization objective. While this is close to our objective, they used fixed penalty parameters, and did not develop the methodology. Besides, [5] proprosed a special bayesian formulation with Gamma prior distributions that is good for the signal problems addressed in the paper but not generically applicable. Moreover, similar to a fused lasso setup, a L1 penalty to smoothen the dynamically changing system together with L1 regularization for sparsity was also proposed by them and [15]. However, these methods are difficult to optimize and, do not guarantee the covergence from an erroneous initialization to a steady-state estimation error.

In contrast to the existing methods, the proposed IRS can easily handle linear and nonlinear (with linear link function) models in high-dimensional space. As we will show in this paper, it provides an efficient sequential modeling approach to accurate model such problems.

3. INERTIAL REGULARIZATION AND SELECTION (IRS)

3.1 Definition and Notations

The following defines an evolutionary model for a data streaming process.

$$\theta_t = F_t\theta_{t-1} + \nu_t \quad (1)$$
$$y_t = X_t\theta_t + \epsilon_t \quad (2)$$

Here t denotes a time epoch for modeling. An epoch can be a day, a week or any other time span, during which the model state is assumed to not change. The observed data at t is denoted by X_t for all independent variables (also referred to as predictors) and y_t is the response. Eq. 2 is conventional statistical model where the response is a linear function of predictors. θ_t is the model parameter and ϵ_t is an error term from a normal distribution with mean 0 and covariance W_t, i.e. $\epsilon_t \sim N(0, W_t)$. Besides, θ_t also specifies the model state at epoch t.

While Eq. 2 models the response, Eq. 1 captures the any change in the state. Often these state changes are small and/or smooth (gradual). As also mentioned before, these changes are called as evolution. Eq. 1 also has a linear structure, where F_t is called state transition parameter and ν_t is the state change process noise, assumed to be $\sim N(0, Q_t)$. Besides, in the rest of the paper we refer Eq. 1 as state-model and Eq. 2 as response-model.

The data sample size and the number of predictor variables at t is denoted as n_t and p_t, respectively. Thus, $X_t \in \mathbb{R}^{n_t \times p_t}$ and $y_t \in \mathbb{R}^{n_t}$. Additionally, the covariance of θ_t is denoted as Σ_t.

3.2 IRS Model

The objective of IRS model is to sequentially learn a statistical model from a streaming data from a high dimensional process. For that, the following objective function is proposed that, upon minimization, yields an optimal model estimate at any epoch t.

$$\mathcal{L}(\theta_t) = \frac{1}{2n_t}(y_t - X_t\theta_t)^T W_t^{-1}(y_t - X_t\theta_t) + \frac{\tau}{2p_t}(\theta_t - F_t\theta_{t-1})^T\Sigma_{t|t-1}^{-1}(\theta_t - F_t\theta_{t-1}) + \frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\theta_{t_i}^*|} \quad (3)$$

The above objective comprises of three components:

a) The first term is a sum of residual squares scaled by the residual error covariance,

b) The second term is a L_2-norm regularization that penalizes any drastic change in the model estimate. A change here implies the difference in the model state at t from the expected model state at t, given the state at $t-1$ ($= F_t\theta_{t-1}$). The second term is, thus, called an inertia component because it resists drastic changes, and in effect, keeps the updated model estimate close to the anticipated model derived from the past model. Besides, this inertia also resists a change due to any noisy observations. These properties are a result of using $\Sigma_{t|t-1}^{-1}$ as an adaptive penalty for an effective model updation, and also helps the inertia component differentiate an abnormal observation from noise (discussed in §3.3 and 6).

c) The third term is a L_1-norm penalty term for an optimal model selection. Although, the L_1 penalty is imposed "spatially", i.e. only on the model parameters at epoch t, the selected model will still resemble with the selected model in the past. This is in part due to the inertial component and an adaptive penalty parameter θ_t^*. This aspect is further elaborated in §3.3.

In Eq. 3, θ_{t-1} is unknown, but it is shown in Appendix-A that it suffices to use the previous epoch's model estimate instead of their true values. This is a founding property for IRS, that allows a sequential (online) modeling. θ_{t-1} can be, thus, replaced with its estimate $\hat{\theta}_{t-1}$. For a succinct notation, we denote $\hat{\theta}_{t|t-1} = F_t\hat{\theta}_{t-1}$, where $\hat{\theta}_{t|t-1}$ is the estimate model state at t given θ_{t-1}. Similarly, the expected covariance of $\hat{\theta}_{t|t-1}$, denoted by $\hat{\Sigma}_{t|t-1}$, is derived from the previous state's covariance, where $\hat{\Sigma}_{t|t-1} = F_t\hat{\Sigma}_{t-1}F_t^T + Q_t$. Likewise, θ_t^* is replaced by $\hat{\theta}_t^*$. The resultant online objective function for sequential modeling is, thus,

$$\hat{\theta}_t = \arg\min_{\theta_t} \frac{1}{2n_t}(y_t - X_t\theta_t)^T W_t^{-1}(y_t - X_t\theta_t) + \frac{\tau}{2p_t}(\theta_t - \hat{\theta}_{t|t-1})^T\hat{\Sigma}_{t|t-1}^{-1}(\theta_t - \hat{\theta}_{t|t-1}) + \frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}_{t_i}^*|} \quad (4)$$

Besides, hyperparameters λ and τ regulates the amount of penalty imposed for limiting the model size (sparsity) and closeness to the expected model given the past, respectively. λ and τ are later shown to be independent of epoch t unless there is any true drastic change in the system. This is an

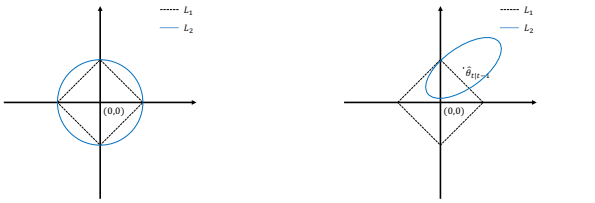

(a) Elastic-net

(b) Inertial Regularization and Selection (IRS)

Figure 1: Graphical illustration of proposed Inertial Regularization and Selection (IRS), and contrasting it with Elastic-net.

important property because it removes the computational overhead of searching optimal λ and $\tau$for every epoch. Besides, each term in the objective is normalized by their size, n_t or p_t, to meet this end.

On another note, Eq. 4 bears close resemblance with an Elastic-net formulation (Zou and Hastie, 2005). Their difference, illustrated graphically in Fig. 1a-1b for a simple 2-dimensional model, further explains and highlights the differentiating properties of IRS. As shown in the figures, the L_2 penalty term in Elastic-net is centered around 0, while IRS's L_2 regularization is around, $\hat{\theta}_{t|t-1}$, the model estimate given the past or the prior knowledge. Moreover, IRS attempts to find a sparse model that also matches the prior knowledge. This reinstates the previous claim that despite of a spatial L_1 penalty the resultant model resembles the model derived from the past.

3.3 Why $\Sigma_{t|t-1}$ and θ^*_t adapted regularization?

In a sequential regularization model, it is critical to update the regularization penalties in a data driven manner. Otherwise, a repetitive regularization optimization for each epoch can be computationally cumbersome. We, therefore, use estimates of $\Sigma_{t|t-1}$ and θ^*_t, which are derived from the data and have the necessary properties, discussed below, that results into an accurate model estimation.

For the inertia part, we want close to the expected value of θ_t given $\hat{\theta}_{t-1}$, i.e. $\hat{\theta}_{t|t-1}$. But the extent of regularization should be different for different θ_{t_i}'s. The θ_{t_i} with higher expected variance should be regularized lesser. This is because, higher variance indicates that the θ_{t_i}'s estimate from the past is not reliable. Thus, it should be learnt more from a recent data from epoch t than keeping close to the past. And it is vice-versa if the expected variance of θ_{t_i} is small — tend to keep it close to a reliable past knowledge. Therefore, the inertia should be weighted by the inverse of the variance. Besides, as the correlations between the model parameters should be taken into account, as well, to adjust the changes in correlated parameters, we use $\Sigma^{-1}_{t|t-1}$ as an adaptive inertia weight.

In addition to the above intuition, $\Sigma^{-1}_{t|t-1}$ also stabilizes the inertial component. The inertial component can be regarded as a sum of residual squares from Eq. 1, i.e. we aim to minimize the residuals of $(\theta_t - \hat{\theta}_{t|t-1})$. Thus, the residual sum of squares should be scaled by its covariance, where $\Sigma_{t|t-1} = \text{cov}(\theta_t - \hat{\theta}_{t|t-1})$.

Besides, an adaptive penalty is applied on the L_1 norm for model selection. As suggested in [18], inverse of a reasonable estimate for θ provides a good adaptive penalty term for L_1 norm. [18] suggested an OLS estimate for a Lasso problem as a good estimate for the adaptive penalty, as it brings estimation consistency and guarantees oracle properties. As shown later, in §3.4, the IRS objective function in Eq. **4** is equivalent to a Lasso problem. From this Lasso equivalence, it can be seen that the OLS estimate desired for adaptively penalizing θ_t is same as an only inertia regularized estimate for θ_t , denoted as $\hat{\theta}^*_t$. The expression for $\hat{\theta}^*_t$ is derived and given in Appendix-B. We, therefore, use $\lambda/|\hat{\theta}^*_{t_i}|$ as penalty for the θ_{t_i}'s L_1 regularization.

3.4 Lasso equivalence

Here we show a Lasso equivalent formulation of IRS. This is an important property because it shows IRS can inherit from a well developed pool of Lasso based techniques. In fact, the use of θ^*_t as the adaptive L_1 penalty for model selection is drawn from this Lasso equivalence. Besides, the estimation algorithm for IRS developed in §3.8 is also inspired from Lasso solutions.

The following shows the Lasso equivalence using a transformation and augmentation of data X_t and y_t.

$$\tilde{X}_t = \begin{bmatrix} \frac{1}{\sqrt{2n_t}} W_t^{-\frac{1}{2}} X_t \\ \sqrt{\frac{\tau}{2p_t}} \Sigma^{-\frac{1}{2}}_{t|t-1} \end{bmatrix} ; \; \tilde{y}_t = \begin{bmatrix} \frac{1}{\sqrt{2n_t}} W_t^{-\frac{1}{2}} y_t \\ \sqrt{\frac{\tau}{2p_t}} \Sigma^{-1/2}_{t|t-1} \hat{\theta}_{t|t-1} \end{bmatrix} \quad (5)$$

Using transformed data in Eq. 5, we can rewrite the objective function in Eq. 4 as a Lasso problem, as given below (see Appendix-C for proof).

$$\mathcal{L}_{\hat{\theta}_{t-1}}(\theta_t) = \left(\tilde{y}_t - \tilde{X}_t\theta_t\right)^T \left(\tilde{y}_t - \tilde{X}_t\theta_t\right) + \frac{\lambda}{p_t} \sum_{i=1}^{p_t} \frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \quad (6)$$

3.5 Kalman Filter as a special case

The Kalman Filter is found to be a special case of IRS. This finding highlights the commonality between filtering methods and regularization.

When $\lambda = 0$, IRS becomes equivalent to a more general Kalman filter which has a weight adjusted state update, i.e. a weighted average of the state estimate from past and a Kalman estimate from the current data. The following shows the IRS estimate expression as equal to a Kalman filter expression (see Appendix-D for the proof).

$$\begin{aligned} \hat{\theta}^*_t &= \left[X_t^T W_t^{-1} X_t + \tau^* \Sigma^{-1}_{t|t-1}\right]^{-1} \cdot \\ &\quad \left[X_t^T W_t^{-1} y_t + \tau^* \Sigma^{-1}_{t|t-1} \hat{\theta}_{t|t-1}\right] \\ &= \hat{\theta}_{t|t-1} + \frac{1}{\tau^*} K_t (y_t - X_t \hat{\theta}_{t|t-1}) \end{aligned} \quad (7)$$

where, $K_t = \Sigma_{t|t-1} X_t^T \left(W_t + \frac{1}{\tau^*} X_t \Sigma_{t|t-1} X_t^T\right)^{-1}$, also known as Kalman gain, and $\tau^* = \tau n_t / p_t$. As we can see above (Eq. 7), the expression for θ_t estimate is same as a conventional (unweighted) Kalman filter when $\tau^* = 1$, and otherwise a weighted Kalman filter when τ^* is any other positive value. Besides, when $\tau^* = 1$, the covariance of $\hat{\theta}^*_t$ can be shown to be same as in Kalman filter.

$$\begin{aligned}\Sigma_t^* &= \text{cov}\left\{\left[X_t^T W_t^{-1} X_t + \Sigma_{t|t-1}^{-1}\right]^{-1} \cdot \right. \\ &\quad \left.\left[X_t^T W_t^{-1} y_t + \Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\right]\right\} \\ &= (I - K_t X_t)\Sigma_{t|t-1} \qquad (8)\end{aligned}$$

It can be noted that, for estimating θ_t and Σ_t, the Kalman filter expression requires inverse of K_t. K_t, however, is a $n_t \times n_t$ matrix, inverting which can be computationally expensive when data size, n_t, is large. We provide an implementation for Kalman filter in Appendix-D, which is significantly faster for a usual scenario of $p_t < n_t$.

3.6 Equivalent Bayesian Formulation

The Bayesian formulation of IRS will have the following prior for θ,

$$\begin{aligned}p_{\lambda,\tau}(\theta_t) \propto \exp - &\left(\frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|}\right) \cdot \\ &\exp\left(-\frac{\tau}{2p_t}(\theta_t - \hat{\theta}_{t|t-1})^T\Sigma_{t|t-1}(\theta_t - \hat{\theta}_{t|t-1})\right) \qquad (9)\end{aligned}$$

This prior is a joint distribution of a double-exponential distribution for the L_1 penalty and a Gaussian centered around $\hat{\theta}_{t|t-1}$ and scaled with the covariance $\Sigma_{t|t-1}$ for the inertial L_2 penalty. The parameters θ_t can be estimated as the Bayes posterior mode (see Appendix-A).

3.7 Model covariance

Due to an L_1 norm term in the loss function, it is difficult to obtain an accurate estimate for the model covariance, Σ_t. Bootstraping can be a numerical way to estimate it. However, bootstrapping can be time consuming. Therefore, we propose an approximate closed form estimate, similar to the one proposed in [16]. In [16], the L_1 penalty $\sum|\theta_i|$ was replaced with $\sum\theta_i^2/|\theta_i|$, and the covariance matrix was approximated at the lasso estimate, $\hat{\theta}$, as a ridge regression expression.

This estimation, however, gives a 0 variance estimate if $\hat{\theta}_i$'s equal to 0. This is not desired in IRS, because a variance estimate for each θ_i is necessary for the inertial component. We, therefore, provide a slightly modified covariance expression, where $|\theta_{t_i}|$ is replaced by $\theta_{t_i}^2/|\hat{\theta}^*_{t_i}|$. Thus, an approximate covariance estimate is given as,

$$\Sigma_t = A_t^{-1}\left(X_t^T W_t^{-1} X_t + \tau^{*2}\Sigma_{t|t-1}^{-1}\right)A_t^{-1} \qquad (10)$$

where,

$$\begin{aligned}A_t &= X_t^T W_t^{-1} X_t + \lambda D_t^{-1} + \tau^*\Sigma_{t|t-1}^{-1} \qquad (11)\\ D_t &= \text{diag}\left(I(\hat{\theta}_{t_i} \neq 0)|\hat{\theta}_{t_i}||\hat{\theta}^*_{t_i}| + \right. \\ &\quad \left.(1 - I(\hat{\theta}_{t_i} \neq 0))|\hat{\theta}^*_{t_i}|^2;\, i = 1, \ldots, p_t\right) \qquad (12)\end{aligned}$$

and $\tau^* = \tau n_t/p_t$.

3.8 Model Estimation

In this section, we first develop a closed form solution under orthogonality conditions, and then present an iterative proximal gradient estimation method for any general case. An algorithm is then laid out for implementing IRS in §3.8.3, where the various computational and implementation aspects are discussed.

3.8.1 Closed-form solution under orthogonality

The developed estimation technique possesses a closed-form solution if, (a) the predictors are orthonormal ($X_t^T X_t = I$), i.e. all columns in X_t are linearly independent of each other, (b) the error of response model is i.i.d., $W_t = w_t^2 I$, and, (c) the error covariance of the state model is i.i.d., $Q_t = q_t^2 I$, which is often implied from if (a) is true.

Orthogonality conditions, a) $X_t^T X_t = I$, b) $W_t = w_t^2 I$, and c) $Q_t = q_t^2 I$. Under these conditions, Σ_t, and in turn $\Sigma_{t|t-1}$, will be diagonal matrices. Suppose we denote, $\Sigma_{t|t-1} = \boldsymbol{\rho}_t I$, where ρ_{t_i} is the expected variance of θ_{t_i}.

Plugging these conditions into Eq. 22, we get the OLS estimate expression as,

$$\begin{aligned}\hat{\theta}_t^* &= \left(\frac{1}{w_t^2}I + \tau^*(\boldsymbol{\rho}_t I)^{-1}\right)^{-1}\left(\frac{1}{w_t^2}X_t^T y_t + \tau^*(\boldsymbol{\rho}_t I)^{-1}\hat{\theta}_{t|t-1}\right) \\ &= (\boldsymbol{\rho}_t^* I)\left(\frac{1}{w_t^2}X_t^T y_t + \tau^*(\boldsymbol{\rho}_t I)^{-1}\hat{\theta}_{t|t-1}\right) \qquad (13)\end{aligned}$$

where, $\tau^* = \frac{\tau n_t}{p_t}$ and $\frac{1}{\rho^*_{t_i}} = \frac{1}{w_t^2} + \frac{\tau^*}{\rho_{t_i}}, i = 1, \ldots, p_t$.

We will use this expression to derive the IRS closed-form solution under orthogonality. From the Lasso equivalent formulation of IRS, the parameter estimate can be computed by minimizing, Eq. 6, i.e.,

$$\mathcal{L}(\theta_t) = \left(\tilde{y}_t - \tilde{X}_t\theta_t\right)^T\left(\tilde{y}_t - \tilde{X}_t\theta_t\right) + \frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \qquad (14)$$

$$\begin{aligned}\arg\min_{\theta_t} \quad & -2\tilde{y}_t^T\tilde{X}_t\theta_t + \theta_t^T\tilde{X}_t^T\tilde{X}_t\theta_t + \frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \\ = \quad & -2\left(\frac{1}{2n_t}y_t^T W_t^{-1}X_t\theta_t + \frac{\tau}{2p_t}\hat{\theta}_{t|t-1}^T\Sigma_{t|t-1}^{-1}\theta_t\right) + \\ & \frac{1}{2n_t}\theta_t^T X_t^T W_t^{-1}X_t\theta_t + \frac{\tau}{2p_t}\theta_t^T\Sigma_{t|t-1}^{-1}\theta_t + \frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \\ \propto \quad & -\left(\frac{1}{w_t^2}y_t^T X_t\theta_t + \tau^*(\boldsymbol{\rho}_t I)^{-1}\hat{\theta}_{t|t-1}^T\theta_t\right) + \\ & \frac{1}{2w_t^2}\theta_t^T\theta_t + \tau^*(\boldsymbol{\rho}_t I)^{-1}\theta_t^T\theta_t + \frac{\lambda n_t}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \\ = \quad & -\left(\frac{1}{w_t^2}y_t^T X_t + \tau^*(\boldsymbol{\rho}_t I)^{-1}\hat{\theta}_{t|t-1}^T\right)\theta_t + \\ & \frac{1}{2}\left(\frac{1}{w_t^2} + \tau^*(\boldsymbol{\rho}_t I)^{-1}\right)\theta_t^T\theta_t + \frac{\lambda n_t}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \\ = \quad & -(\boldsymbol{\rho}_t^* I)^{-1}\hat{\theta}_t^{*T}\theta_t + \frac{1}{2}(\boldsymbol{\rho}_t^* I)^{-1}\theta_t^T\theta_t + \frac{\lambda n_t}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \\ = \quad & \sum_{i=1}^{p_t} -\frac{1}{\rho^*_{t_i}}\hat{\theta}^*_{t_i}\theta_{t_i} + \frac{1}{2\rho^*_{t_i}}\theta_{t_i}^2 + \frac{\lambda n_t}{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|} \qquad (15)\end{aligned}$$

\

On solving this for both cases when, $\theta_i^{(t)} \geq 0$ and ≤ 0, we get the following solution,

$$\hat{\theta}_{t_i} \quad = \quad \text{sgn}(\hat{\theta}^*_{t_i}) \left(\left| \hat{\theta}^*_{t_i} \right| - \frac{\lambda n_t}{p_t} \frac{\rho^*_{t_i}}{|\hat{\theta}^*_{t_i}|} \right)^+ \tag{16}$$

where, $(\cdot)^+$ is a soft-thresholding function.

The components in the solution expression (Eq. 16) can be interpreted as a shrinkage on the inertia regularized parameter (OLS) estimate, $\hat{\theta}^*_{t_i}$, such that the amount of shrinkage is determined by, a) the magnitude of $\hat{\theta}^*_{t_i}$, for an adaptive variable selection from L_1 penalty, and b) $\rho^*_{t_i}$, which is proportional to the expected variance of the OLS estimate, $\hat{\theta}^*_{t_i}$, thus, IRS strongly shrinks $\hat{\theta}^*_{t_i}$ if its expected variance is high.

Therefore, parameters with high variance and low magnitude will be removed. The addition of variance element boosts the model selection capability to select a better model, resulting into a lower prediction errors.

3.8.2 Iterative estimation

Since, most commonly found problems do not satisfy an orthogonality condition, we develop a proximal gradient method to estimate the parameters in a general case.

$$\begin{aligned} \mathcal{L}(\theta_t) \quad &= \quad \underbrace{\left\{ \frac{1}{2n_t} (y_t - X_t\theta_t)^T W_t^{-1} (y_t - X_t\theta_t) + \frac{\tau}{2p_t} (\theta_t - \hat{\theta}_{t|t-1})^T \Sigma^{-1}_{t|t-1} (\theta_t - \hat{\theta}_{t|t-1}) \right\}}_{f(\theta_t) \rightarrow \text{differentiable}} \\ &\quad + \frac{\lambda}{p_t} \sum_{i=1}^{p_t} \underbrace{\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|}}_{g(\theta_t) \rightarrow \text{non-differentiable}} \\ &= \quad f(\theta_t) + g(\theta_t) \end{aligned}$$

Since, g is a L_1 function, an iterative proximal gradient algorithm will perform the following update,

$$\theta^{l+1}_{t_i} \quad \leftarrow \quad \text{Prox}_\nu \left(\theta^l_{t_i} - s^l \nabla f(\theta^l_{t_i}) \right) \tag{17}$$

where, $\nu = \frac{\lambda s^l}{p_t |\hat{\theta}^*_{t_i}|}$ and $\text{Prox}_\nu(x) = \begin{cases} x - \nu & ; x \geq \nu \\ 0 & ; |x| < \nu \\ x + \nu & ; x \leq -\nu \end{cases}$, s^l is a step size for iteration l (the step size can be kept same for all l or decreased with increasing l) and

$$\nabla f(\theta^l_t) = -\frac{1}{n_t} (X_t)^T W_t^{-1} (y_t - X_t\theta^l_t) + \frac{\tau}{p_t} \Sigma^{-1}_{t|t-1} (\theta^l_t - \hat{\theta}_{t|t-1}) \tag{18}$$

Note that, any other estimation technique, for example, subgradient method or Least-Angle Regression (LARS), for estimating θ. We chose proximal gradient method because of its ease of understanding and implementation. Besides, it performs well in most circumstances.

3.8.3 Algorithm

In this section, we present Algorithm-1 for implementing the iterative estimation method discussed above. The algorithm takes in the predictor matrices, X_t, and the responses, y_t, as they are collected for epoch $t = 1, 2, \ldots$. The X_t's are standardized within each epoch t, such that the columns in X_t has a mean of 0 and variance of 1. y_t, on the other hand, is centered to remove an intercept from the model.

The algorithm assumes F_t and Q_t as known inputs. In practice, F_t is often set as equal to an identity matrix, $F_t = I$, and Q_t as $\acute{o}^2 I$, where $\acute{o}$ is a small positive number. Besides, the response error, ϵ_t, in Eq. 2, is assumed as i.i.d., i.e. $W_t = w^2 I$, and $p_t = p, \forall t$. While these are reasonable assumptions for most applications, it can be easily extended for circumstances outside of these assumptions.

As shown in Algorithm-1, the IRS estimation process starts from line-24. To initialize, an OLS estimate is computed, and residual sum of squares is used for obtaining w^2. Besides, the covariance is initialized as identity. Thereafter, for any t, $\hat{\theta}_{t|t-1}$, $\Sigma_{t|t-1}$ and w^2 (for approximating W_t) are computed and passed into ProximalDescent function, alongwith the data and regularization parameters, (λ, τ). This function implements the proximal gradient steps shown above in §3.8.2. Note that, it takes inverse of $\Sigma_{t|t-1}$ to save computation. Besides, it is recommended to have relatively larger data size in the initial epoch to avoid yielding a suboptimal initial model.

A common issue in regularization models is tuning the penalty parameters, viz. (λ, τ) in IRS. It would have been even more challenging if IRS required to (re)tune (λ, τ) for each estimation epoch. However, due to the adaptive regularization developed in §3.3, it can be shown that the expected value of the loss function, $\mathcal{L}_{\hat{\theta}_{t-1}}(\theta_t)$, remains a constant, i.e. independent of epoch t (see Appendix **??**). This implies that an appropriately chosen (λ, τ) can be used for any $t = 1, 2, \ldots$.

Therefore, for IRS, we can use a k-fold cross validation (CV) on a two-dimensional grid search for (λ, τ), and model first few epochs for each grid point and measure their prediction errors. We used a 10-fold CV on first three epochs for (λ, τ) selection. Zou and Hastie (2005) have described this approach in detail and provided other possible selection approaches.

Computational time is another common issue in estimation procedures. The developed proximal gradient based algorithm has a fast convergence, which can be further boosted by changing the stepsize, s, for each iteration in line 11-18. This is common practice, with one simple approach as starting with a large value for s and gradually decreasing it after each iteration. Another approach is to adaptively increase or decrease s based on an increase or decrease in the residual error from θ^l_t. Using this, the proximal gradient was found to converge in less than 50 iterations for upto 1000-dimensional models.

Besides, in implementation, operations in lines 13-16 can be vectorized and approximate matrix inversion methods can be used for a faster computation. Furthermore, accelerated proximal gradient and stochastic proximal gradient methods can be used to boost the computation.

3.9 Ability to handle missing values and structural changes

A beneficial feature of IRS is the ease of handling missing values and any change in model structure. A change in model structure implies any addition or removal of model predictors. In case of missing values, IRS can still provide a reasonable parameter estimate in a modeling epoch by using the prior knowledge.

If a new predictor variable is added, for eg., one new product at Walmart, it can be easily incorporated in the modeling by expanding $\hat{\theta}_{t|t-1}$ and $\hat{\Sigma}_{t|t-1}$. If there is no expert prior

Algorithm 1 Inertial Regularization and Selection (IRS) estimation

Input: $X_t, y_t, \hat{\theta}_{t-1}, \Sigma_{t-1}, \lambda, \tau, Q_t; t = 1, 2, \ldots$

function PROX(x, ν)

$$x \leftarrow \begin{cases} x - \nu & x > \nu \\ 0 & |x| \leq \nu \\ x + \nu & x < -\nu \end{cases}$$

return x

end function

function PROXIMALDESCENT$(X_t, y_t, \hat{\theta}_{t|t-1}, \Sigma^{-1}_{t|t-1}, w^2, \lambda, \tau)$

$n_t \leftarrow \arg_{n_t} X_t \in \mathbb{R}^{n_t \times p}$

$\tau^* \leftarrow \tau n_t / p$

$W_t \leftarrow w^2 I_{n_t}$

$\hat{\theta}^*_t \leftarrow (X_t^T W_t^{-1} X_t + \tau^* \Sigma^{-1}_{t|t-1})^{-1}(X_t^T W_t^{-1} y_t + \tau^* \Sigma^{-1}_{t|t-1} \hat{\theta}_{t|t-1})$

Initialize:

$l = 0, \theta^0_t = \hat{\theta}^*_t, s$

while $l < L$ **do**

$\nabla f(\theta^l_t) \leftarrow \frac{-1}{n_t} X_t^T W_t^{-1}(Y_t - X_t \theta^l_t) + \frac{\tau}{p} \Sigma^{-1}_{t|t-1}(\theta^l_t - \hat{\theta}_{t|t-1})$

for $i \in 1, \ldots, p$ **do**

$\nu \leftarrow \frac{\lambda s}{p|\hat{\theta}^*_{t_i}|}$

$\theta^l_{t_i} \leftarrow$ PROX$(\theta^l_{t_i} - s \nabla f(\theta^l_t)_i, \nu)$

end for

$l \leftarrow l + 1$

end while

$D_t \leftarrow \text{Diag}(\mathbf{1}(\hat{\theta}_{t_i} \neq 0)|\hat{\theta}_{t_i}||\hat{\theta}^*_{t_i}| + (1 - \mathbf{1}(\hat{\theta}_{t_i} \neq 0))|\hat{\theta}^*_{t_i}|^2; i = 1, \ldots, p)$

$A_t \leftarrow X_t^T W_t^{-1} X_t + \lambda D_t^{-1} + \tau^* \Sigma^{-1}_{t|t-1}$

$\Sigma_t \leftarrow A_t^{-1}(X_t^T W_t^{-1} X_t + \tau^{*2} \Sigma^{-1}_{t|t-1}) A_t^{-1}$

return $\hat{\theta}_t, \Sigma_t$

end function

Initialize:

$n_1 \leftarrow \arg_{n_1} X_1 \in \mathbb{R}^{n_1 \times p}$

$\hat{\theta}_0 \leftarrow (X_1^T X_1)^{-1} X_1 y_1$

$w^2 \leftarrow (y_1 - X_1 \hat{\theta}_0)^T (y_1 - X_1 \hat{\theta}_0)/(n_1 - 1)$

$\Sigma_0 \leftarrow I_p$

IRS Estimation for, $t = 1, 2, \ldots$

$n_t \leftarrow \arg_{n_t} X_t \in \mathbb{R}^{n_t \times p}$

$\hat{\theta}_{t|t-1} \leftarrow F_t \hat{\theta}_{t-1}$

$\Sigma_{t|t-1} \leftarrow F_t \Sigma_{t-1} F_t^T + Q_t$

$w^2 \leftarrow (y_t - X_t \hat{\theta}_{t|t-1})^T (y_t - X_t \hat{\theta}_{t|t-1})/(n_t - 1)$

$\hat{\theta}_t, \Sigma_t \leftarrow$

PROXIMALDESCENT$(X_t, y_t, \hat{\theta}_{t|t-1}, \Sigma^{-1}_{t|t-1}, w^2, \lambda, \tau)$

Output Model estimate: $\hat{\theta}_t, \Sigma_t$

knowledge for the new predictor, $\hat{\theta}_{t|t-1}$ can be expanded with 0 values and $\hat{\Sigma}_{t|t-1}$ with 0 off-diagonal (no correlation) and a large variance. This is same as an uninformative prior used is a Bayesian estimation. However, in several cases there is some expert knowledge. For instance, in the above Walmart example if the new product is a new baby diaper, we will have some knowledge for its parameter from similar products, and/or its correlations with other baby and related products.

Besides, IRS can handle removal of a predictor more cleverly than most other conventional methods. A removed predictor does not mean its effect will be gone immediately. In fact it will take some time before its effect vanishes completely. For eg. suppose a TV brand, say XYZ, is removed from a Walmart store, but since XYZ's sales must be interacting with other TV sales, the effect of its absence will be felt for some time (maybe raising the demand for others) before gradually dying out. In IRS, it is not necessary to manually remove any predictor variable from the model even if we know it is physically removed. The inertial and selection component in IRS will automatically update the resultant effect and eventually remove the predictor variables from the model when its effect declines completely.

It is worthwhile to note that this IRS property also handles situations when the removal of the predictor is temporary. In the above Walmart eg., a reintroduction of the TV XYZ can be easily incorporated in the modeling updates, with the effect of reintroduction being the opposite of removal — the effect gradually increasing in the model.

This ability to easily handle missing values and structural changes in model is a valuable property because it makes the sequential modeling more robust to the real world scenarios.

4. EXPERIMENTAL VALIDATION

In this section, we experimentally validate the efficacy of IRS. We compare it with Lasso, Kalman Filter (KF) and Ensemble KF (EnKF). Lasso was applied locally on each epoch $t$to show a prediction error baseline if the prior knowledge is ignored in sequential process. KF and EnKF are commonly used state-of-the-art state-space models against which IRS's performance will be gauged.

We setup two types of experiments for validation, denoted as Exp-1 and Exp-2. In Exp-1, the number of model parameters (predictors) are kept same for all epochs, $p_t = p = \{500, 5000\}, t = 1, 2, \ldots, T$. The model parameters are randomly initialized for $t = 1$ from a multivariate normal distribution with mean as 0, and an arbitrary covariance, such that a sparsity level of 0.2 is maintained. As the model enters from epoch t to $t + 1$, its parameters change with a random walk of variance 1. The sample size is randomly selected between $(1.8p, 2.1p)$. The data is generated for 9 epochs, i.e. $T = 9$. A 10-fold CV is performed on initial three epochs sample for regularization parameter selection. Thereafter, again a 10-fold CV is performed to fit each method, and the average test prediction error, rMSEs, are shown in Fig. 2a-2b.

The figures show that Lasso has a similar rMSE as IRS in the initial epoch. Thereafter, while Lasso's prediction error remain approximately at the same level, IRS's error reduces. This is due to the use of knowledge used from the past for updating the model with current data. KF and EnKF similar had a reducing error rates, however, they had

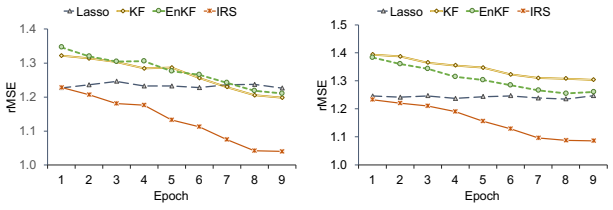

(a) $p = 500$ (b) $p = 5000$

Figure 2: Exp-1: Average rMSE of test errors from a 10-fold cross validation.

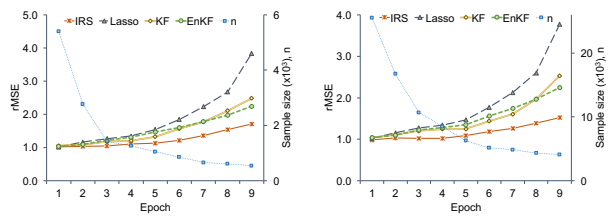

(a) $p = 500$ (b) $p = 5000$

Figure 3: Exp-2: Average rMSE of test errors from a 10-fold cross validation.

a higher error than IRS and Lasso for the most of epochs. This can be attributed to their inability to work in a high-dimensional sparse processes. Besides, their performance seems to worsen when $p = 5000$, compared to when $p = 500$. This maybe due to KF and EnKF overfitting the data, resulting into higher test prediction errors, that gets higher as dimension increases.

In Exp-2, we initialized the model similar as Exp-1. However, for subsequent epochs, $t > 1$, a parameter in θ_t was changed as, a) if θ_{t-1_i} is 0, it can become non-zero at t with a small activation probability ($= 0.05$), and its value is again drawn from the initial multivariate normal distribution, b) if $\theta_{t-1_i} \neq 0$, then it can be deactivated (made equal to 0) if it is less than a deactivation threshold ($= 0.1$) with a probability ($= 0.1$), c) if $\theta_{t-1_i} \neq 0$ and it is not deactivated, then it is changed with a random walk, $F_{t_i} = 1$, or directionally $F_{t_i} = (1 + N(0, 0.1))$ with added random noise ($\sim N(0, 1)$).

This setup emulates a real world problem where, an otherwise insignificant variable becomes significant and vice-versa, due to any systemic changes. Also, the shift in θ is sometimes directional, indicating an evolution. Regardless, we will perform the sequential modeling assuming $F_t = I$, to stay in tandem with the common state-space modeling practice.

In addition to incorporating evolution, we also reduce the data sample size for progressing epochs to test each methods' efficacy in working with growing prior knowledge and lesser data.

Again, a 10-fold CV is performed and the average rMSEs from test prediction is reported in Fig. 3a-3b. As expected, the errors increase with passing epochs due to lesser sample size. While in the initial epochs, all methods have roughly the same error rate, it increases sharply for a localized Lasso fit. Between, EnKF and KF, the latters error rate grows faster and surpasses the rMSE of EnKF after some epochs pass. This can be due to EnKF's better ability to estimate the parameter covariance even with lesser data (using the Monte Carlo setup). Among all, the increase in error rate is the lowest for IRS.

This result shows that IRS is also robust to the evolutionary model changes, and can work with lesser data as it garners more prior knowledge. This is an important feature for sequential modeling — indicating the ability of quickly updating the model as the system gets older.

The results in this section show the importance of use of past information for sequentially modeling data that arrives in a streaming fashion. For high-dimensional sparse and evolving models, it is shown that IRS outperforms the state-space models popular for such sequential modeling. Next, we will show a real world application of IRS.

5. APPLICATION ON AN ONLINE RETAIL DATA

Here we show a real world application of IRS using a retail data from a UK-based and registered non-store online retail. The company mainly sells unique all-occasion gifts to mostly wholesalers and some individuals. The company was established in 1981 and relied on direct mailing catalogues. It went online in about 2010, since when it started accumulating huge amount of data.

A subset of that data used in this study lies between Dec, 2010 to Dec, 2011. The data is a customer transaction dataset, which has several variables, namely, product name, quantity sold, invoice date-time, unit price, customer-id, and country. For our analysis, we drop the variable customer-id, and select data from the country, UK. From this filtered data, we take a subset belonging to 41 products, such that it comprises of high to low selling items.

The "quantity sold" for a product in a transaction is the response variable, and the model's objective is to predict it using all available predictors. For predictors, we extract the day-of-week and quarter-of-day from the invoice date-time to form predictors. Month is not taken because the data is for just one year, thus, deriving month seasonal effects will not be possible. Besides, other predictors are the unit price and product name. We create dummy columns for the categorical product name and day-of-week variables. Finally, a predictor matrix is made by taking in all predictor variables and their second-order interactions. Interaction effects are important to consider because sales of one product usually effects sales of some others. In total, the predictor matrix has 363 columns, that is, we have a 363-dimensional model.

Each month is treated as a modeling epoch. A 10-fold CV is performed on the first three epochs for penalty parameters, (λ, τ), selection, and then IRS is performed sequentially on all epochs. The 10-fold CV test prediction error (mean absolute percentage error, MAPE) in Fig. 4 show that IRS is more accurate than others. This is because several predictors in the above model will be statistically insignificant, thus, would require an effective model selection, in addition to use of the prior knowledge. However, as also mentioned before, KF and EnKF cannot effectively find a parsimonious model, and, thus, has poorer test accuracy.

In addition, we show another result of prediction errors for low selling products in Fig. 5. Demand forecast of low selling items are usually more difficult because of fewer available data. As shown in the figure, IRS again outperfoms other

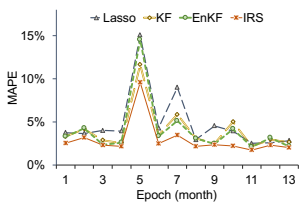

Figure 4: Forecasting accuracy for a randomly selected set of items.

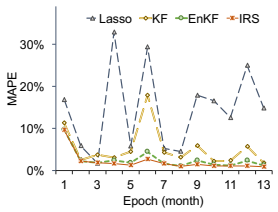

Figure 5: Forecasting accuracy for low selling items.

models. This can be attributed to IRS's better capability is working with lesser data.

An online retail data, such as the one used here, is often very high-dimensional due to several categorical variables and interactions. Besides, the demand patterns typically have an inherent evolutionary characteristics. Note that, this evolution is different from the seasonalities and trend. Evolution is a gradual change due to various external factors, that are hard to measure. IRS is, thus, well suited for such problems. It can account for the evolutionary changes, and can handle high-dimensions, as well as, an automatic model selection.

6. DISCUSSION AND CONCLUSION

In the previous section (§5), we show IRS's superior prediction accuracy in an online retail data problem. This is an example of a high-dimensional statistical learning problem with sparsity and large data collection in a streaming fashion, where the process can also evolve. A majority of statistical learning problems found these days have similar properties. For instance, sensors data, telecommunications, weblogs, console logs, user music/movie play history (on services like spotify and Netflix), etc. Besides, in such problems various variables get added or removed from the process over time. Under all this fast-paced dynamics, IRS can be easily applied and yield a high accuracy.

In general, IRS's ability to auto-tune its regularization parameters (given in §3.3) from the data makes it easy to implement. Due to this auto-tune ability, as further elucidated in §3.8.3, IRS has just two hyperparameters, (λ, τ), which if suitably selected provides an optimal solution for any epoch $t = 1, 2, \ldots$. However, in event of any abnormal or extreme system change, (λ, τ) may need to be re-tuned. Nevertheless, under normal conditions IRS provides a seamless sequential estimation for a statistical learning model.

Besides, the combination of inertial and L_1 penalty terms bring important properties of *resemblance* and *model selection*. Due to the inertia penalty, the model does not drastically change between epochs, thus, keeps a resemblance with what an analyst already knows about the model. The L_1 penalty provides an appropriate model selection on top of it. It must be noted that in model selection, the prior knowledge is also incorporated by using an adaptive penalty. This prevents the L_1 penalty from selecting an extremely different model than the past.

However, a potential downside of the *resemblance* property could be if the initial model is severely sub-optimal, the subsequent models can also be poor. In such cases, it may take several epochs for the model estimate to get close to the optimal. Therefore, to avoid this, it is recommended in §3.8.3 to use a larger amount of data in the initial epoch.

Regardless, IRS is also robust to noise and missing values due to the inertia component. This, alongwith an optimal model selection via L_1 penalty significantly boosts the test prediction accuracy, demonstrated in §4-5.

In this paper, we show IRS's equivalence with Lasso, Kalman Filter and Bayesian estimation in §3.4-3.6. This also shows the interconnection between them for a statistical learning. Specifically, Kalman filters are often seen as a different family of estimators isolated from regularization techniques. Here, we establish their connection, and also present a faster implementation for Kalman filter in Appendix-D.

In summary, we develop a new method – Inertial Regularization and Selection (IRS) – for sequentially modelin an evolving statistical learning model. We present the modeling approach for a linear state-change (evolution) and a linear state-response model with gaussian errors. However, it can be easily extended to any non-linear model, if there is a link function transforming it into a linear form. For eg., a binary response in a classification model can be modeled using a logit-link for the state response model. The method is shown to outperform existing methods and several real world applications.

7. REFERENCES

[1] E. Candes and J. Romberg. Sparsity and incoherence in compressive sampling. *Inverse problems*, 23(3):969, 2007.

[2] E. J. Candes, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted $\ell 1$ minimization. *Journal of Fourier analysis and applications*, 14(5-6):877–905, 2008.

[3] A. Carmi, P. Gurfil, and D. Kanevsky. Methods for sparse signal recovery using kalman filtering with embedded pseudo-measurement norms and quasi-norms. *IEEE Transactions on Signal Processing*, 58(4):2405–2409, 2010.

[4] A. Charles, M. S. Asif, J. Romberg, and C. Rozell. Sparsity penalties in dynamical system estimation. In *Information Sciences and Systems (CISS), 2011 45th Annual Conference on*, pages 1–6. IEEE, 2011.

[5] A. S. Charles and C. J. Rozell. Dynamic filtering of sparse signals using reweighted ℓ 1. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6451–6455. IEEE, 2013.

[6] Y. Ding and X. Li. Time weight collaborative filtering. In *Proceedings of the 14th ACM international conference on Information and knowledge management*, pages 485–492. ACM, 2005.

[7] G. Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367, 2003.

[8] N. J. Gordon, D. J. Salmond, and A. F. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE Proceedings F-Radar and Signal Processing*, volume 140, pages 107–113. IET, 1993.

[9] S. J. Julier and J. K. Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422, 2004.

[10] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.

[11] J. Lei and P. Bickel. Ensemble filtering for high dimensional non-linear state space models. *University of California, Berkeley, Rep*, 779:23, 2009.

[12] J. S. Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.

[13] E. Ott, B. R. Hunt, I. Szunyogh, A. V. Zimin, E. J. Kostelich, M. Corazza, E. Kalnay, D. Patil, and J. A. Yorke. A local ensemble kalman filter for atmospheric data assimilation. *Tellus A*, 56(5):415–428, 2004.

[14] R. L. Plackett. Some theorems in least squares. *Biometrika*, 37(1/2):149–157, 1950.

[15] D. Sejdinovic, C. Andrieu, and R. Piechocki. Bayesian sequential compressed sensing in sparse dynamical systems. In *Allerton Conf. Communication, Control, and Computing*. Citeseer, 2010.

[16] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[17] N. Vaswani. Kalman filtered compressed sensing. In *2008 15th IEEE International Conference on Image Processing*, pages 893–896. IEEE, 2008.

[18] H. Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.

APPENDIX

A. ONLINE RECURSIVE ESTIMATION LOSS FUNCTION

The parameter estimate is a Bayesian formulation can be found from the Maximum a Posteriori (MAP) function, as shown below,

$$\begin{aligned}\{\hat{\theta}_t\}_{t=\{1,\dots T\}} &= \arg\max \left[\left(\prod_{t=1}^{T} p(\theta_t|\theta_{t-1}) p(y_t|\theta_t \right) \right. \\ &\quad \left. p(y_0|\theta_0) p(\theta_0) \right] \\ &= \arg\max \left[p(y_T|\theta_T) p(\theta_T|\theta_{T-1}) \cdot \right. \\ &\quad \left(\prod_{t=1}^{T-1} p(\theta_t|\theta_{t-1}) p(y_t|\theta_t) \right) \\ &\quad \left. p(y_0|\theta_0) p(\theta_0) \right] \end{aligned} \tag{19}$$

A marginalization can be performed on Eq. 19 to find a globally optimal solution at the T^{th} epoch.

$$\begin{aligned}\hat{\theta}_T &= \arg\max_{\theta_T} \left[\int_{\mathbb{R}^T} p(y_T|\theta_T) p(\theta_T|\theta_{T-1}) \left(\prod_{t=1}^{T-1} p(\theta_t|\theta_{t-1}) \cdot \right. \right. \\ &\quad \left. \left. p(y_t|\theta_t) \right) \cdot p(y_0|\theta_0) p(\theta_0) d\{\theta_i\}_{i=1,\dots,T-1} \right] \\ &= \arg\max \left[p(y_T|\theta_T) p(\theta_T|\theta_{T-1}) \int_{\mathbb{R}^T} \left(\prod_{t=1}^{T-1} p(\theta_t|\theta_{t-1}) \cdot \right. \right. \\ &\quad \left. \left. p(y_t|\theta_t) \right) p(y_0|\theta_0) p(\theta_0) d\{\theta_i\}_{i=1,\dots,T-1} \right] \end{aligned} \tag{20}$$

The integral part in Eq. 20 is essentially the prior on θ_T. Thus, for an online estimation, the updated form of Eq. 19 for θ_T, given below, can be used.

$$\hat{\theta}_T = \arg\max_{\theta_T} p(y_T|\theta_T) p_{\hat{\theta}_{T-1}}(\theta_T) \tag{21}$$

Eq. 21 summarizes all the past information from the likelihood $\theta_T|\{y_t\}_{t=1,\dots,T-1}$ into the prior for θ_T given $\hat{\theta}_{T-1}$. Using the Bayesian prior distribution for IRS given in §3.6, into Eq. 21, we get,

$$\begin{aligned}\hat{\theta}_T &= \arg\max_{\theta_T} \exp\left[-(y_T - X_T\theta_T)^T W_T^{-1} (y_T - X_T\theta_T) \right] \cdot \\ &\quad \exp\left[-\tau(\theta_T - \hat{\theta}_{T|T-1})^T \Sigma_{T|T-1}^{-1} (\theta_T - \hat{\theta}_{T|T-1}) - \right. \\ &\quad \left. \sum \lambda_{t_i} \frac{|\theta_{T_i}|}{|\hat{\theta}^*_{T_i}|} \right] \\ &= \arg\min_{\theta_T} (y_T - X_T\theta_T)^T W_T^{-1} (y_T - X_T\theta_T) + \\ &\quad \tau(\theta_T - \hat{\theta}_{T|T-1})^T \Sigma_{T|T-1}^{-1} (\theta_T - \hat{\theta}_{T|T-1}) + \sum \lambda_{t_i} \frac{|\theta_{T_i}|}{|\hat{\theta}^*_{T_i}|} \end{aligned}$$

B. OLS ESTIMATE FOR THE LASSO EQUIVALENT

The OLS estimate, denoted by $\hat{\theta}^*_t$, for IRS can be easily found from its lasso equivalent given in Eq. 6, in §3.4.

$$\begin{aligned}\hat{\theta}^*_t &= (\tilde{X}_t^T \tilde{X}_t)^{-1} \tilde{X}_t^T \tilde{y}_t \\ &= \left(\frac{1}{2n_t} X_t^T W_t^{-1} X_t + \frac{\tau}{2p_t} \Sigma_{t|t-1}^{-1} \right)^{-1} \cdot \\ &\quad \left(\frac{1}{2n_t} X_t^T W_t^{-1} y_t + \frac{\tau}{2p_t} \Sigma_{t|t-1}^{-1} \hat{\theta}_{t|t-1} \right) \\ &= \left(X_t^T W_t^{-1} X_t + \tau^* \Sigma_{t|t-1}^{-1} \right)^{-1} \left(X_t^T W_t^{-1} y_t + \tau^* \Sigma_{t|t-1}^{-1} \hat{\theta}_{t|t-1} \right) \end{aligned} \tag{22}$$

where, $\tau^* = \frac{\tau n_t}{p_t}$.

C. PROVING LASSO EQUIVALENCE

Using an augmented X and y in Eq. 5, we need to show,

$$\begin{aligned}&(\tilde{y}_t - \tilde{X}_t\theta_t)^T (\tilde{y}_t - \tilde{X}_t\theta_t)_{\text{L.H.S.}} \\ &= \left[\frac{1}{2n_t} (Y_t - X_t\theta_t)^T W_t^{-1} (Y_t - X_t\theta_t) + \right. \\ &\quad \left. \frac{\tau}{2p_t} (\theta_t - \hat{\theta}_{t|t-1})^T \Sigma_{t|t-1}^{-1} (\theta_t - \hat{\theta}_{t|t-1}) \right]_{\text{R.H.S.}} \end{aligned}$$

To prove L.H.S. is equal to R.H.S., we expand L.H.S.:

$$
\begin{aligned}
& (\tilde{y}_t - \tilde{X}_t\theta_t)^T(\tilde{y}_t - \tilde{X}_t\theta_t) \\
= \;& \tilde{y}_t^T\tilde{y}_t - 2\tilde{y}_t^T\tilde{X}_t\theta_t + \theta_t^T\tilde{X}_t^T\tilde{X}_t\theta_t \\
= \;& \frac{1}{2n_t}y_t^TW_t^{-1}y_t + \frac{\tau}{2p_t}\hat{\theta}_{t|t-1}^T\Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1} \\
& -2\left(\frac{1}{2n_t}y_t^TW_t^{-1}X_t\theta_t + \frac{\tau}{2p_t}\hat{\theta}_{t|t-1}^T\Sigma_{t|t-1}^{-1}\theta_t\right) \\
& +\frac{1}{2n_t}\theta_t^TX_t^TW_t^{-1}X_t\theta_t + \frac{\tau}{2p_t}\theta_t^T\Sigma_{t|t-1}^{-1}\theta_t \\
= \;& \frac{1}{2n_t}\left(y_t^TW_t^{-1}y_t - 2y_t^TW_t^{-1}X_t\theta_t + \theta_t^TX_t^TW_t^{-1}X_t\theta_t\right) \\
& +\frac{\tau}{2p_t}\left(\theta_t^T\Sigma_{t|t-1}^{-1}\theta_t - 2\hat{\theta}_{t|t-1}^T\Sigma_{t|t-1}^{-1}\theta_t + \hat{\theta}_{t|t-1}^T\Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\right) \\
= \;& \frac{1}{2n_t}(y_t - X_t\theta_t)^TW_t^{-1}(y_t - X_t\theta_t) + \\
& \frac{\tau}{2p_t}(\theta_t - \hat{\theta}_{t|t-1})^T\Sigma_{t|t-1}^{-1}(\theta_t - \hat{\theta}_{t|t-1}) \\
= \;& \text{R.H.S.}
\end{aligned}
$$

D. SHOWING KALMAN FILTER EQUIVALENCE

Here we expand the IRS expression for θ_t estimate when $\lambda = 0$, denoted by $\hat{\theta}_t^*$, given in Eq. 22 in Appendix-B.

$$
\begin{aligned}
\hat{\theta}_t^* &= \left[X_t^TW_t^{-1}X_t + \tau^*\Sigma_{t|t-1}^{-1}\right]^{-1}\cdot \\
&\quad \left[X_t^TW_t^{-1}y_t + \tau^*\Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\right] \\
&= \left[\frac{1}{\tau^*}\Sigma_{t|t-1} - \frac{1}{\tau^*}K_tX_t\frac{1}{\tau^*}\Sigma_{t|t-1}\right]\cdot \\
&\quad \left[X_t^TW_t^{-1}y_t + \tau^*\Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\right] \\
&= \hat{\theta}_{t|t-1} - \frac{1}{\tau^*}K_tX_t\hat{\theta}_{t|t-1} + \\
&\quad \left[\frac{1}{\tau^*}\Sigma_{t|t-1}X_t^TW_t^{-1} - \frac{1}{(\tau^*)^2}K_tX_t\Sigma_{t|t-1}X_t^TW_t^{-1}\right]y_t \\
&= \hat{\theta}_{t|t-1} - \frac{1}{\tau^*}K_tX_t\hat{\theta}_{t|t-1} + \\
&\quad K_t\left[\frac{1}{\tau^*}\left(W_t + \frac{1}{\tau^*}X_t\Sigma_{t|t-1}X_t^T\right)W_t^{-1} - \right. \\
&\quad \left.\frac{1}{(\tau^*)^2}X_t\Sigma_{t|t-1}X_t^TW_t^{-1}\right]y_t \\
&= \hat{\theta}_{t|t-1} + \frac{1}{\tau^*}K_t(y_t - X_t\hat{\theta}_{t|t-1}) \qquad (23)
\end{aligned}
$$

where, $K_t = \Sigma_{t|t-1}X_t^T\left(W_t + \frac{1}{\tau^*}X_t\Sigma_{t|t-1}X_t^T\right)^{-1}$, also known as Kalman gain.

Besides, when $\tau^* = 1$, the covariance of $\hat{\theta}_t^*$ can be shown to be same as in Kalman filter.

$$
\begin{aligned}
\Sigma_t^* &= \text{cov}\left\{\underbrace{\left[X_t^TW_t^{-1}X_t + \Sigma_{t|t-1}^{-1}\right]^{-1}}_{E_t}\cdot\right. \\
&\quad \left.\left[X_t^TW_t^{-1}y_t + \Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\right]\right\} \\
&= E_t^{-1}X_t^TW_t^{-1}\text{cov}(y_t)W_t^{-1}X_t^TE_t^{-1} + \\
&\quad E_t^{-1}\Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\Sigma_{t|t-1}^{-1}E_t^{-1} \\
&= E_t^{-1}\left(X_t^TW_t^{-1}X_t^T + \Sigma_{t|t-1}^{-1}\right)E_t^{-1} \\
&= E_t^{-1}E_tE_t^{-1} \\
&= \left(X_t^TW_t^{-1}X_t^T + \Sigma_{t|t-1}^{-1}\right)^{-1} \\
&= \Sigma_{t|t-1} - \\
&\quad \Sigma_{t|t-1}X_t^T\left(W_t + X_t\Sigma_{t|t-1}X_t^T\right)^{-1}X_t\Sigma_{t|t-1} \\
&= \Sigma_{t|t-1} - K_tX_t\Sigma_{t|t-1} \\
&= (I - K_tX_t)\Sigma_{t|t-1}
\end{aligned}
$$

Faster Kalman filter algorithm

Given a previous estimate of θ and its covariance, as $\hat{\theta}_{t-1}$ and Σ_{t-1}, and the data, the updated model is found using the following steps,

$$
\begin{aligned}
\hat{\theta}_{t|t-1} &= F_t\hat{\theta}_{t-1} && (24)\\
\Sigma_{t|t-1} &= F_t\Sigma_{t-1}F_t^T + Q_t && (25)\\
K_t &= \Sigma_{t|t-1}X_t^T\left(\Sigma_t + X_t\Sigma_{t|t-1}X_t^T\right)^{-1} && (26)\\
\hat{\theta}_t &= \hat{\theta}_{t|t-1} + K_t\left(y_t - X_t\hat{\theta}_{t|t-1}\right) && (27)\\
\Sigma_t &= \left(I - K_tX_t\right)\Sigma_{t|t-1} && (28)
\end{aligned}
$$

The formulation shows an evolving model as an additive function of the prior model. However, computationally it becomes intensive if the datasize is even reasonably large. This is due to a $n_t \times n_t$ matrix inversion required in Eq. 26. Using the IRS equivalence for Kalman, we can have the following procedure for Kalman filter estimation, which is significantly faster when $p < n_t$, i.e. the number of parameters being less than the datasize, which is more common in statistical learning models (unlike the signal processing model, a primary focus for Kalman filters).

A faster algorithm for Kalman will, thus, replace Eq. 26-28 as,

$$
\begin{aligned}
\hat{\theta}_t &= \left[X_t^TR_t^{-1}X_t + \Sigma_{t|t-1}^{-1}\right]^{-1}\left[X_t^TR_t^{-1}y_t + \Sigma_{t|t-1}^{-1}\hat{\theta}_{t|t-1}\right] \\
\Sigma_t &= \left(X_t^TR_t^{-1}X_t + \Sigma_{t|t-1}^{-1}\right)^{-1}
\end{aligned}
$$

E. STATIONARITY OF IRS LOSS FUNCTION

The IRS loss function given in Eq. 4 has three components. Here we will show that they are individually stationary, meaning they have a constant expectation, independent of time epochs.

$$\mathcal{L}_{\hat{\theta}_{t-1}}(\theta_t) = \underbrace{\frac{1}{2n_t}(y_t - X_t\theta_t)^T W_t^{-1}(y_t - X_t\theta_t)}_{\text{(A)}} + \underbrace{\frac{\tau}{2p_t}(\theta_t - \hat{\theta}_{t|t-1})^T \Sigma_{t|t-1}^{-1}(\theta_t - \hat{\theta}_{t|t-1})}_{\text{(B)}} + \underbrace{\frac{\lambda}{p_t}\sum_{i=1}^{p_t}\frac{|\theta_{t_i}|}{|\hat{\theta}^*_{t_i}|}}_{\text{(C)}} \quad (29)$$

We will use the following identity applicable for a random variable, $x \sim N(\mu, \Sigma)$, and any known matrix A.

$$E[x^T Ax] = \text{Tr}(A\Sigma) + \mu^T A\mu \quad (30)$$

- Part (A): $y_t - X_t\theta_t \sim N(0, W_t)$, where W_t is a symmetric matrix. Therefore, $E[(\text{A})] = \frac{1}{2n_t}\left(\text{Tr}(W_t^{-1}W_t) + 0\right) = \frac{1}{2n_t}n_t = \frac{1}{2}$.
- Part (B): $\theta_t - \hat{\theta}_{t|t-1} \sim N(0, \Sigma_{t|t-1})$. Similarly, as for part (A), we will have, $E[(\text{B})] = \frac{\tau}{2}$.
- Part (C): Suppose, θ_{t_i} is approximated as a hard-thresholded estimate from OLS, $\theta_{t_i} = \hat{\theta}^*_{t_i} \cdot \epsilon^*_{t_i}$, where $\epsilon^*_{t_i} \sim \text{Bernoulli}(\alpha)$. Therefore, $E\left[\left|\frac{\theta_{t_i}}{\hat{\theta}^*_{t_i}}\right|\right] = \alpha \implies E[(\text{C})] = \lambda\alpha$. Here, α can be interpreted as the proportion of variables that are selected from the L_1 penalization. Assuming α is same for all epochs, the expectation of part (C) is independent of t.